Integrated Information and Network Management for End-to-End Quality of Service

Marco Carvalho and
Adrian Granados
Florida Institute for Human and
Machine Cognition - Ocala, FL
{mcarvalho, agranados}@ihmc.us

Kyle Usbeck, Joseph Loyall
and Matthew Gillen
BBN Technologies
Cambridge, MA
{kusbeck, jloyall, mgillen}@bbn.com

Asher Sinclair and
James P. Hanna
Air Force Research Laboratory
Rome, NY
{Asher.Sinclair, James.Hanna}@rl.af.mil

Abstract—**Publish-subscribe-based Information Management (IM) services provide a key enabling technology for net-centric operations. This paper describes technology for Quality of Service (QoS) and Internet-Protocol-based Airborne Networking features for IM services. Enhancing IM services with airborne networking features improves effectiveness in combined tactical and enterprise networks with mobile airborne and ground-based embedded platforms interacting with enterprise systems in command and control operations.**

Keywords-component; Airborne Networks, QoS Management, Cross-Layer Network Design, QED, Cross-Layer Substrate

I. INTRODUCTION

Information Management (IM) services provide key technology for net-centric operations by enabling the active discovery, management, and dissemination of information based on content and metadata, through a publish/subscribe/query paradigm that decouples information producers and consumers. IM services in mission-critical systems must include the capability to specify and enforce Quality of Service (QoS) requirements for information objects, which must be satisfied under the constraints of the available communications and computational resources.

In contrast, Network Management (NM) refers to the monitoring and control of networking parameters to ensure message delivery properties are within cost and capacity constraints. NM is information agnostic; focuses on lower-level metrics and operates at the physical, data link, network and transport layers; and is generally designed and optimized for specific types of networks and topologies.

The increasing interconnectivity, complexity, and heterogeneity of deployed communication systems and the combination of enterprise and tactical networks provide challenges for effective network and QoS management for multiple users and potentially competing missions. Enterprise networks typically run Internet or intranet protocols, relying on high bandwidth and relatively stable communication links and Internet Protocol (IP)-based communication. In contrast, tactical networks frequently service mobile nodes running Mobile Ad hoc Network (MANET) communication protocols, and have both constrained and dynamic network conditions.

In this paper we introduce an integrated information and network management approach for QoS support in mixed tactical and enterprise networks consisting of mobile nodes with ad hoc communication moving between terrestrial networks utilizing traditional IP-based protocols. In these environments, QoS features common in enterprise networks, such as Differentiated Services (DiffServ) [2], do not readily translate to IM service-level QoS in tactical networks due to the high potential for lossy, noisy, and intermittent links. For example, the loss of only a single packet at the network layer might statistically be considered high QoS, but might prevent a message from being reconstructed at the IM service-level, rendering the other packets delivered for that message useless.

Our approach builds upon previous work in QoS management for IM services under the QoS Enabled Dissemination (QED) project [10] and a cross-layer communication substrate [3] that provides airborne networking capabilities for addressing of mobile nodes, bandwidth and bottleneck awareness, and prioritization of information at the application- and packet-level.

In this paper we present our design and implementation of QoS management for airborne networks, a prototype implementation for a US Air Force-developed set of IM services [7], and a corresponding demonstration and experimentation results.

II. FEATURES OF AIRBORNE NETWORKS

Airborne Networks (ANs) are characterized as a combination of fixed and ad hoc network components. Airborne platforms communicate through wireless links to one another generally using ad hoc protocols, and to nodes in terrestrial networks generally using static routes, or infrastructure-based routing protocols, as shown in Figure 1.

The scenario illustrated in Figure 1 requires a number of key capabilities that include addressing of mobile nodes as they migrate between networks, prioritized data delivery, link capacity estimation, and bandwidth management on shared wireless links. We discuss each of these in turn.

Report Documentation Page

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
NOV 2011		**00-00-2011 to 00-00-2011**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Integrated Information and Network Management for End-to-End Quality of Service	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Raytheon BBN Technologies,Cambridge,MA,02138**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
MILCOM, Baltimore, MD, November 7-10, 2011. U.S. Government or Federal Rights License

14. ABSTRACT
Publish-subscribe-based Information Management (IM) services provide a key enabling technology for net-centric operations. This paper describes technology for Quality of Service (QoS) and Internet-Protocol-based Airborne Networking features for IM services. Enhancing IM services with airborne networking features improves effectiveness in combined tactical and enterprise networks with mobile airborne and ground-based embedded platforms interacting with enterprise systems in command and control operations.

15. SUBJECT TERMS

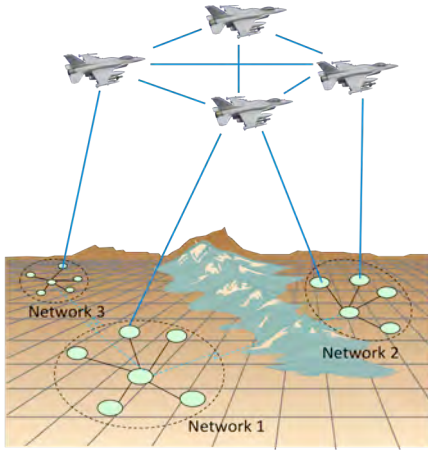| 16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**6**	

Figure 1. Interpretation of an airborne network.

A. Mobile IP

Mobile IP [13] enables aerial nodes to seamlessly transition between ground networks and continue communicating, without the need for explicit reconfiguration – a key capability for airborne networks. Mobile IP allows *mobile* nodes to maintain their IP addresses when moving across separate networks. Hence, ongoing sessions (e.g., TCP streams) can be maintained during network migrations.

B. Differentiated Services

The purpose of Differentiated Services (DiffServ) is to allocate bandwidth through prioritized treatment of various classes of traffic. A Weighted Fair Queuing (WFQ) approach to DiffServ ensures that each traffic flow has fair access to network resources, i.e., no traffic flows are capable of starving other flows and bursty flows are prevented from consuming more than their share of bandwidth. WFQ, however, does not allocate dedicated amounts of bandwidth per-priority, as this would result in sub-optimal use of the network. Instead, if bandwidth is available, it is consumed by the highest priority traffic available at the time-of-send. DiffServ can be provided at multiple OSI layers (e.g., application layer, network layer, etc.).

C. Bandwidth Estimation and Bottleneck Detection

QoS management is facilitated by an accurate estimation of the capacity of the network. This is especially useful to determine how much information of various priorities the available network capacity can support. Otherwise, it is possible for an IM layer to introduce more high priority information into the network than it can support, resulting in the network level dropping or in the delay of individual packets, each of which can prevent the reconstruction of an important message.

Accurate bandwidth estimation has been difficult to do reliably in practice. Fixed terrestrial networks have been able to rely on theoretical capacities – reliable because of the fixed nature of the network. In contrast, ANs have network capacities that vary significantly due to distance, noise, and other factors, rendering any theoretical capacities of the network unreliable.

A useful approximation is afforded by *bottleneck detection*. Since an AN is made up of a network of individual platform to platform links, detecting a particular link that has become a

bottleneck (i.e., more traffic has arrived at the link than the link can transmit at any given time) allows QoS management to reduce the load caused by messages traversing that link, even if the actual capacity of the network is unknown.

D. Explicit Channel Reservation

Explicit Channel Reservation (ECR) provides guaranteed bandwidth for a stream of important packets. ECR works by establishing reservations for a stream of traffic. Although DiffServ is more prevalent, protocols like the ReSerVation Protocol (RSVP) [17] are still used occasionally for important message delivery or applications with streaming data.

E. Concurrent Multipath Routing

The purpose of Concurrent Multipath Routing (CMR) is to increase network utilization by exploiting multiple network paths between nodes. CMR can be used to increase bandwidth between nodes and/or balance load by routing packets over different network paths. CMR can also be used for fault tolerance by providing redundant paths for network packets.

III. IMPLEMENTING AN FEATURES IN THE VIA COMMUNICATIONS SUBSTRATE

VIA is a next generation cross-layer communications substrate [3] for tactical networks and IM systems. VIA enables applications to adapt and leverage the characteristics of the dynamic communication environment and enables the underlying communications infrastructure to better support application QoS requirements and constraints.

As shown in Figure 2, VIA operates below the network layer, which allows it to transparently manage protocol packets such as ARP and IP, including ICMP, TCP and UDP. VIA does not require encapsulation and control information is inserted into IP packets, e.g., using the IP Options field. Thus, existing applications can take advantage of VIA without any modifications and non-VIA and VIA nodes can interoperate.
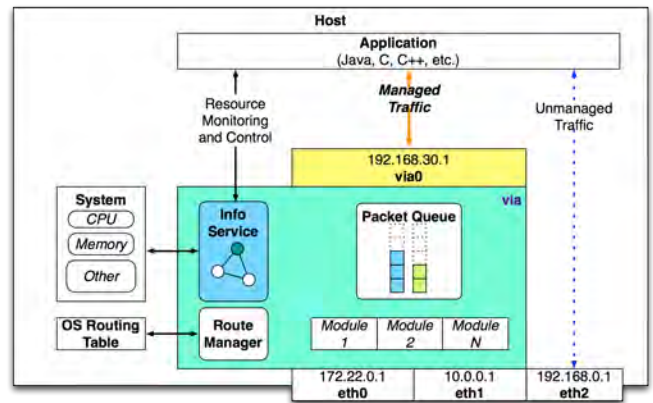

Figure 2. Conceptual view of the VIA architecture.

Applications that bind to VIA's virtual network interface (i.e., *via0*) can transparently make use of VIA capabilities to support communication and QoS requirements. Outgoing packets are intercepted, processed, and sent over one or more physical interfaces operating in promiscuous mode. Conversely, incoming packets are captured, processed, and injected into the virtual network interface, emulating a physical link between the two end-points.

Additional features that extend existing VIA capabilities are implemented as *modules*. Based on the capabilities described in Section II, we identified a set of capabilities to be created as VIA modules to support AN features. Each of these implemented modules, including Mobile IP, Bottleneck Detection, Differentiated Services, and Capacity Estimation are described in this section.

A. Mobile IP

As discussed in the previous sections, Mobile IP allows *mobile* nodes to maintain their IP addresses when moving across separate networks, maintaining active streams during network migrations. Assuming that at least one node in each network can communicate with one another using normal IP routing mechanisms, transparent migration of *mobile* nodes can be achieved using a combination of ad hoc routing and transparent packet tunneling.

Consequently, VIA implements an ad hoc based approach to Mobile IP that provides a mechanism to propagate ad hoc routes across the mobile node's *home* and *foreign* networks[1], as well as the ability to automatically establish a Generic Routing Encapsulation (GRE) tunnel between two *gateway* nodes that transparently forwards traffic from the home to the foreign network and vice-versa. Upon joining a foreign network, mobile nodes register themselves with the gateway node (Figure 3). When the mobile node migrates to a foreign network (1), it registers with the foreign gateway node (2), which exchanges registration information with the mobile node's home gateway node to create and maintain the GRE tunnel (3) for enabling the propagation of packets. After migration, VIA-enabled nodes in the home and foreign networks automatically update the routes to the mobile node and communicate through the GRE tunnel transparently.
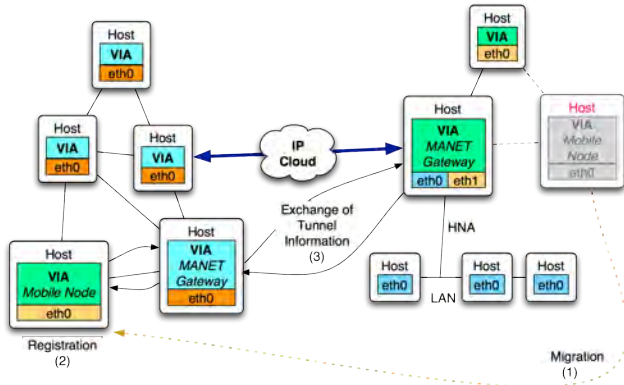

Figure 3. A mobile node migrates from its home network to a foreign network.

B. Bottleneck Detection

Communication paths between nodes in airborne and terrestrial networks consist of shared bandwidth links, as well as links with different (possibly dynamic) bandwidth capacities that could result in bottlenecks when transmitting data at high rates. Moreover, intermediate nodes have different resource constraints that may limit the speed at which packets can be forwarded at each hop, causing delays that result in communication bottlenecks. To detect a link bottleneck, VIA keeps track of the number of bytes transmitted (TX rate) and received (RX rate) for all frames, and periodically reports back to the upstream node the actual RX rate. Each upstream node compares the RX rate reported by the downstream node to the actual TX rate, and if the difference between rates is greater than a certain threshold, VIA emits a bottleneck notification (Figure 4). Since the goal is to provide IM service-level QoS on an end-to-end basis, the notification specifies the IP addresses of the affected destinations[2], rather than the next hop IP information, and an estimate of the capacity through the bottleneck.
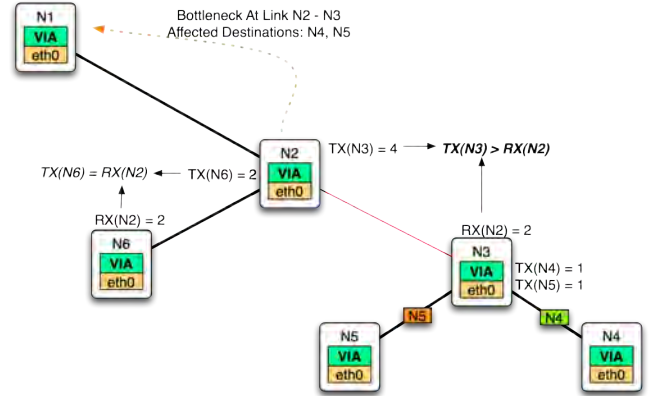

Figure 4. Node N2 detects a bottleneck and reports it back to node N1.

In addition, VIA periodically monitors the size of the transmission queue(s). If the size of the queue(s) is larger than a certain threshold, VIA emits a bottleneck notification. In contrast to what occurs when a link bottleneck is detected, the addresses of the affected destinations are determined by inspecting the IP header of each of the packets waiting in the queue.

C. Differentiated Services

VIA provides support for differentiated services using a WFQ approach. A module classifies packets into eight (0-7) distinct traffic classes based on the type of traffic and, in the case of IP datagrams, on the value of the Differentiated Service Code Point (DSCP) field of the packet's IP header. Each traffic class has a corresponding queue with a pre-configured weight that allows the module to prioritize the transmission of packets and guarantee a minimum bandwidth allocation for each traffic class. With the goal of maintaining network connectivity at all times when possible, the module automatically classifies packets from certain protocol types such as ARP and other network management traffic, including VIA control messages, as traffic class 7, which has the highest priority (i.e., traffic from this class is always sent ahead of any other type of traffic).

D. Capacity Estimation

QoS features, such as bandwidth reservation and rate adaptation, require knowledge of available link capacity to effectively manage the utilization of the link. A naïve method for capacity estimation consists of saturating the link by sending

---

[1] The home network is the network from which the mobile node's persistent IP address is derived. A foreign network is any network into which the mobile node moves.

[2] The addresses of the affected destinations are determined by inspecting the IP header of each packet that traverses the node.

artificial traffic at an increasing transmission rate, until packet loss is detected at the application layer. The main drawback with this approach is the disruption of application traffic that occurs when the link is operating over capacity.

VIA utilizes an approach for on-demand capacity estimation in wireless links with Media Access Control (MAC) that also relies on sending artificial traffic to saturate the link but introduces little to no disruption to the application traffic. It relies on adverse link conditions being manifested at the MAC layer before they affect the behavior of application traffic. VIA monitors the size of the MAC transmission queue to determine the saturation point of the link, looking for the queue to grow above a certain threshold before packet loss would be detected at the application layer. Specifically, VIA monitors the number of frames queued by the MAC layer, and in parallel introduces artificial traffic at a low transmission rate. While the number of MAC frames waiting to be transmitted is below a certain threshold, VIA increases the transmission rate of artificial traffic. The remaining link capacity is assumed to be the reported transmission rate just before the size of the queue reaches the threshold, and at which VIA stops the transmission of artificial traffic. By detecting the queue growth at the MAC layer, VIA is able to compute the capacity without packet loss at the application layer.

IV. INTEGRATED INFORMATION QoS SUPPORT

While QoS management in the airborne network is a necessary feature for control and visualization, it is not sufficient to achieve aggregate QoS control across a mission or to enforce higher-level client or information-level policies, such as treating all packets of an information object uniformly, enforcing deadlines, replacing stale information, or changing format to fit client needs or available bandwidth. What is needed is monitoring and control of the underlying communications architecture by a higher level QoS manager, as shown in Figure 5.
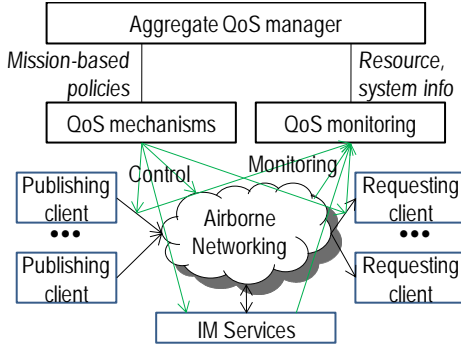
Figure 5. Achieving mission-based QoS requires management and control at the application and networking layers.

To prototype an integrated system for IM service-level QoS management in ANs, we have integrated our QED services [9], [10] with the AN features provided by VIA. QED provides a set of policy-driven QoS managers and mechanisms that ensure the timely and smooth brokering and dissemination of important information through IM services, gracefully degrade when resources are overloaded, and enforce client preferences and priorities. QED includes a Dissemination Service that performs prioritization of information delivery, bandwidth management, and information shaping, and a Task Manager that

schedules CPU intensive tasks. These local QoS managers (local because they manage QoS at a single control point) are under the control of an Aggregate QoS Manager that stores, selects, and distributes policies to achieve end-to-end, consistent QoS for multiple publish-subscribe information flows.

We have previously described QED's architecture, prototype implementation, and mechanisms and services [9], [10], as well as its use in embedded US Air Force, Marine, and Navy exercises [6], [12]. To provide QoS in IM systems, the QED prototype, shown in Figure 6, performs the following:

- Manages the scheduling of threads and bandwidth typically hidden behind service interfaces.
- Provides and enforces QoS policy specified at a high level, based on user, information, and system concepts.
- Provides policy-driven aggregate QoS management across control points and across users, despite the publishing and consuming clients being decoupled from one another.
- Handles information that varies in size and processing that varies in time.
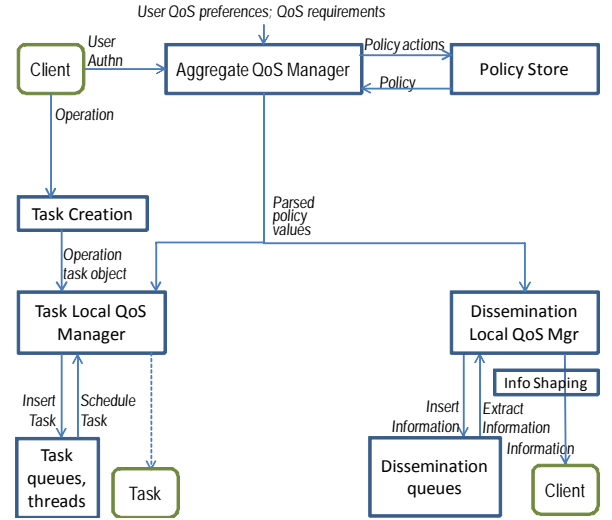
Figure 6. QED provides QoS management for IM Services.

QED prioritizes information objects based on high-level mission policies, enforces deadlines, and replaces old information objects. It utilizes VIA's AN features to maintain, publish and subscribe connections to mobile clients. QED receives bottleneck notifications from VIA and throttles the rate of information to match the reported capacity to clients sharing bottleneck links.

V. EXPERIMENTAL RESULTS

We ran experiments to determine the effectiveness of Mobile IP and Bottleneck Detection in airborne networks. All experiments were run with VIA and QED.

A. Mobile IP Experiment

The goal for the Mobile IP experiment is to determine the amount of time required for an airborne mobile subscriber to reconnect to a terrestrial network after a Mobile IP migration and resume receipt of Information Objects (IOs) from the IM services (without re-subscribing). To measure this value, we perform the following steps: a) Publish constant-size IOs at a

rate that does not exceed network capacity, b) Trigger a mobile IP migration by disconnecting the mobile subscriber from one airborne network and connecting it to the other airborne network, and c) Measure the time between the last IO seen before the migration and the first IO seen after the migration.

B. Mobile IP Results

The results of the Mobile IP experiment are shown in Table 1. The results represent thirty (30) iterations of the mobile IP experiment in each direction of migration.

Table 1. Mean and standard deviation of reconnection times after a mobile IP migration in each migration direction.

	Reconnection Time (sec)	
Migration Direction	**Mean**	**σ**
Foreign to Home	2.3	1.9
Home to Foreign	9.0	4.0

Our first finding was that the reconnection time after a Mobile IP migration depends on the direction of the migration. Migrating from the home network to a foreign network is significantly slower because of the overhead associated with creating a tunnel between the two networks. The tunnel is not required for communicating within the home network, so the reconnection time when migrating back to the home network is significantly faster.

We also found that VIA was able to handle mobile IP migration fast enough that there was often no noticeable disturbance to the application-layer. Migrations, however, can cause noticeable network delays, and thus can cause problems in applications that are not tolerant to network disruptions.

C. Bottleneck Management Experiment

The goal of the bottleneck management experiment is to determine the effect of bottleneck management on prioritized IOs in a bottleneck-constrained network from the point-of-view of the information subscriber. The metric we use to measure bottleneck management's performance is end-to-end IO latency. To measure this value, we performed the following steps:

- Publish IOs at a rate that exceeds network link capacity to the subscriber, but does not exceed the network link capacity from any publisher.
- Measure the time elapsed from the creation of the IO by its producer (the publisher) to the time it is received by its consumer (the subscriber).

D. Bottleneck Management Results

The results of the bottleneck management experiment are shown in Figure 7. These results represent two iterations of the bottleneck management experiment – one with bottleneck management enabled ("With Mgmt"), and one with bottleneck management disabled ("No Mgmt"). The results show that, with bottleneck management enabled, high-priority traffic is strictly preferred over low-priority. It also shows that low-priority traffic does not pay a significant penalty for operating with bottleneck management.

QED tries to send high-priority traffic first, but when the link capacity is unknown, it cannot properly prioritize the traffic. This fact explains the discrepancy between the "High-Priority (With Mgmt)" and "High-Priority (No Mgmt)." The lack of link capacity knowledge causes inconsistency in prioritization and can even result in priority inversion (as seen at the beginning of the experiment). The reason this occurs is that once the application-layer sends an IO to the transport layer, it can no longer control its outbound priority. It becomes just another set of packets in the interface's egress buffer.
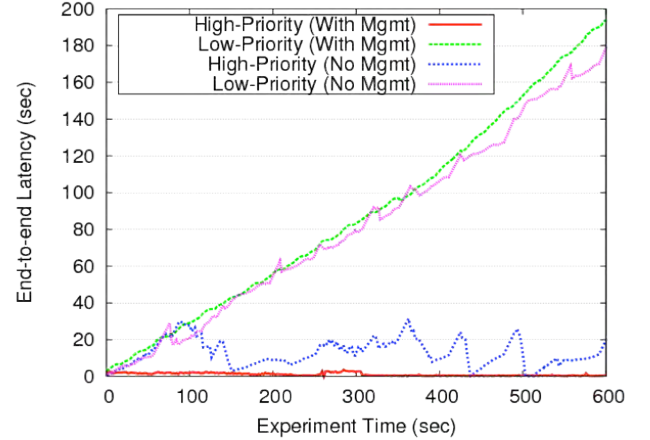

Figure 7. End-to-end latency of high/low priority IOs with bottleneck management enabled versus disabled.

Another problem we encountered occurs when probing for an increase in link capacity. At this point, TCP slow-start was throttling traffic at the transport layer before VIA could detect the bottleneck. To resolve this issue, we replaced the TCP congestion control algorithm with one that does not slow-start. Another solution to the problem would be to use a transport that does not perform congestion control (e.g., UDP).

VI. RELATED WORK

While there has not been a lot of work on combined information and network management approaches for airborne network environments, there has been previous independent research in each of publish-subscribe Information Management and Network Management.

The primary information management research that we build upon is based on the Joint Battlespace Infosphere concept [14], [15], which has been realized in several prototype versions [4], [7], [8], [16]. A number of other publish-subscribe services are available, although most only provide an interface for disseminating information, rather than the rich, active information management of the IM services. Eugster *et al.* provide an overview of the pub-sub interaction pattern, highlighting the decoupled nature of publishers and subscribers in time, space, and synchronization [5]. A few researchers have investigated QoS management in pub-sub middleware. Mahambre *et al.* present a taxonomy of pub-sub middleware with QoS features, acknowledging significant gaps in the provision of QoS features to the extent that some environments need [11].

The Joint Capability for Airborne Networking (JCAN) [1] is a multi-layer infrastructure for airborne network communications that provides different features at the data-link, network

and transport layers. It also includes application layer libraries for interacting with the lower-layer features of JCAN. JCAN provides implementations for Mobile IP, ECR WFQ, and CMR.

The key difference in JCAN's implementation of Mobile IP is the mechanism JCAN uses for intercepting all the network traffic destined for the mobile nodes, called Proxy ARP. With Proxy ARP, the gateway on the home network responds to ARP requests made for any mobile nodes, and forwards the intercepted traffic over an IP tunnel.

JCAN's implementation of ECR has several components. One component is an application layer interface to link reservation. Another component is a replacement C library for Unix sockets, called *gSocket*, for rate control. This socket library has the same structure as the Unix socket library for easy replacement in source code, however, every program that wants to make use of gSockets must be recompiled. Another shortcoming of JCAN's ECR implementation is that a multi-hop link must be fully reserved before any data is transferred. Similarly, reservations hinder network utilization by preventing traffic flows from using a currently under-utilized link. For these reasons, differentiated services have been used more frequently.

JCAN's WFQ module uses an Active Timed Queue (ATQ) Linux kernel module, a replacement to the normal outgoing packet scheduler. The ATQ scheduler is the component that prevents traffic flows from starving each other of bandwidth. The downside of using the ATQ module, as opposed to something that provides "strict" priority, is that the ATQ scheduler relies on accurate link capacity estimation, which is not available in JCAN.

JCAN implements CMR using the *iproute2* suite to manipulate IP-layer routing. The *ip route* command allows multiple *nexthop* arguments, which specify the alternatives for egress routes. JCAN's implementation of CMR uses the Linux kernel *equalize* patch, which allows transport-layer flows (e.g., TCP streams) to be split among multiple egress routes. However, the *equalize* kernel patch has been deprecated since the development of the JCAN framework, and therefore, we were unable to get the patch to work (despite using a JCAN-prescribed version of Linux). Furthermore, the Linux Advanced Networking community recommends against using the patch for various reasons, but mostly because it prevents optimizations that improve the performance of some transport protocols.

VII. CONCLUSIONS

Our initial prototype of Airborne Networking features and integration with IM and QoS management services has shown promise in enabling greater levels of predictability, control, and performance for net-centric information exchange with mobile platforms and heterogeneous wireless and wired networks. Our approach goes beyond previous approaches in supporting (1) a combination of ad hoc and IP-based routing which supports faster routing to locally-discoverable nodes and (2) QoS management for information objects, which provides better mission-oriented quality while achieving high levels of network utilization.

The research and prototype software that we describe establishes a basis upon which to build for more complete AN and QoS features in heterogeneous operations at large scale. Future

work includes establishing more of the AN features such as ECR and CMR, addressing inconsistencies between IP-based and tactical networks, and evaluating the capabilities at larger scale and in embedded, tactical platforms.

REFERENCES

[1] Baskinger, P., Card, S., Chruscicki, M., Demase, M., Smetek, R., and Zabele, S., "Information for Global Reach," US AFRL Technical Report AFRL-IF-RS-TR-2005-308, http://handle.dtic.mil/100.2/ADA439303.

[2] Blake S., Black D., Carlson M., Davies E., Wang Z., and Weiss W., "An Architecture for Differentiated Services," IETF Network Working Group RFC 2475, (1998).

[3] Carvalho, M., Granados A., Naqwi, W., Brothers, A., Hanna, J., and Turck, K., "A cross-layer communications substrate for tactical information management systems," Proc. of the Military Communications Conference (MILCOM), (2008).

[4] Combs, V., Hillman, R., Muccio, M., and McKeel, R., "Joint battlespace infosphere: information management within a C2 enterprise," Proc. Tenth Int. Command and Control Technology Symp. (ICCRTS), (2005).

[5] Eugster, P., Felber, P., Guerraoui, R., and Kermarrec, A.-M., "The many faces of publish/subscribe," ACM Computing Surveys, 35(2):114-131, 2003.

[6] Gillen, M., Loyall, J., and Sterling, J., "Dynamic Quality of Service Management for Multicast Tactical Communications," Proc. 14th IEEE Computer Society Symposium on Object/Component/Service-oriented Real-time Distributed Computing (ISORC), Newport Beach, CA, (2011).

[7] Grant, R., Combs, V., Hanna, J., Lipa, B., and Reilly, J., "Phoenix: SOA based information management services," Proc. of the SPIE Defense Transformation and Net-Centric Systems Conference, Orlando, Fl, (2009).

[8] Linderman, M., Siegel, B., Ouellet, D., Brichacek, J., Haines, S., Chase, G., and O'May, J., "A reference model for information management to support coalition information sharing needs," Proc. Tenth Int. Command and Control Technology Symp. (ICCRTS), (2005).

[9] Loyall, J., Gillen, M., Paulos, A., Edmondson, J., Varshneya, P., Schmidt, D., Bunch, L., Carvalho, M., and Martignoni III, A., "Dynamic policy-driven quality of service in service-oriented systems," Proc. 13th IEEE Computer Society Symposium on Object/component/service-oriented Real-Time distributed Computing (ISORC), Carmona, Spain, (2010).

[10] Loyall, J., Sinclair, A., Gillen, M., Carvalho, M., Bunch, L., Martignoni III, A., and Marcon, M., "Quality of service in US Air Force information management systems," Proc. of the Military Communications Conference (MILCOM), (2009).

[11] Mahambre, S., Kumar, M., and Bellur, U., "A taxonomy of QoS-aware, adaptive event-dissemination middleware," IEEE Internet Computing, 11(4):35-44, 2007.

[12] Paulos, A., Sinclair, A., and Loyall, J., "Evaluating QoS-Enabled Information Management Services in a Navy Operational Context," Proc. SPIE Conference on Defense Transformation and Net-Centric Systems, Orlando, FL (2011).

[13] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, Aug. 2002.

[14] Scientific Advisory Board, "Building the joint battlespace infosphere volume 1: summary," SAB-TR-99-02, (2000).

[15] Scientific Advisory Board, "Report on building the joint battlespace infosphere volume 2: interactive information technologies," SAB-TR-99-02, (1999).

[16] Uschold, M., Clark, P., Dickey, F., Fung, C., Smith, S., Uczekaj, S., Wilke, M., Bechhofer, S., and Horrocks, I., "A Semantic Infosphere," 2nd International Semantic Web Conference (ISWC 2003), Sanibel Island, FL, October 20-23, 2003.

[17] Zhang, L., Deering, S., Estrin, D., Shenker, S., and Zappala, D., "RSVP: A New Resource ReSerVation Protocol," IEEE Network, September 1993.